


NASA/TM—2017-219510



Performance of the Widely-Used CFD code OVERFLOW on the Pleides Supercomputer

*Guru P. Guruswamy
Ames Research Center, Moffett Field, California*

May 2017

NASA STI Program ... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:


- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.
- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at <http://www.sti.nasa.gov>
- E-mail your question to help@sti.nasa.gov
- Fax your question to the NASA STI Information Desk at 443-757-5803
- Phone the NASA STI Information Desk at 443-757-5802
- Write to:
STI Information Desk
NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320

NASA/TM—2017-219510



Performance of the Widely-Used CFD code OVERFLOW on the Pleides Supercomputer

*Guru P. Guruswamy
Ames Research Center, Moffett Field, California*

National Aeronautics and
Space Administration

*Ames Research Center
Moffett Field, CA 94035-1000*

May 2017

Available from:

NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320
443-757-5802

This report is also available in electronic form at
<http://www.nas.nasa.gov/publications/reports/reports.html>

Performance studies of the Widely-Used OVERFLOW CFD Code on the Pleiades Supercomputer

Guru P. Guruswamy

Abstract

Computational performance studies were made for NASA's widely used Computational Fluid Dynamics code OVERFLOW on the Pleiades Supercomputer. Two test cases were considered: a full launch vehicle with a grid of 286 million points and a full rotorcraft model with a grid of 614 million points. Computations using up to 8000 cores were run on Sandy Bridge and Ivy Bridge nodes. Performance was monitored using times reported in the day files from the Portable Batch System utility. Results for two grid topologies are presented and compared in detail. Observations and suggestions for future work are made.

Introduction

Under the NASA's High End Computing Capability (HECC) project, the NASA Advanced Supercomputing (NAS) Division of Ames Research Center is continuously enhancing the capabilities of NASA's supercomputers, including the Pleiades supercomputer system [1,2]. In parallel, continuous improvements are being made in the OVERFLOW code [3], which is widely used (about 100 active users) for many aerospace applications [4, 5]. NAS routinely conducts in depth performance studies on simpler topologies [6]. However it is useful and prudent to periodically assess the performance of OVERFLOW for real world complex geometries on existing nodes and newly installed of Pleiades hardware.

The OVERFLOW code solves the Reynolds-Averaged Navier-Stokes Equations using overset grid topologies generated by grid tools such as OVERGRID [7]. Typically, grids can be grouped into two types: near-body (NB) grid-zones that are coincident with no-slip surfaces and off-body (OB) grid-zones that connect near-body grid zones with free stream. NB grids are typically supplied via an external grid generation process whereas OB grids are either generated through automated grid generation tools in OVERFLOW or supplied by an external grid generator.

In this report parallel performance of version 2.2g of OVERFLOW for two physical cases, each with a different grid topology: a full launch vehicle model of the Saturn V [8] using user-generated OB grids, and a full rotorcraft wind-tunnel model of the HART II [9] with OVERFLOW-generated OB grids. Studies are based Sandy Bridge (San) nodes and Ivy Bridge (Ivy) nodes that were successfully operating during 2010 to 2015. Performances of these two grid topologies were monitored using times reported by OVERFLOW and Portable Batch System (PBS) [1] protocol, a computational job queue system.

Nodes Used

The Pleiades supercomputer system includes 1872 San nodes with 2 eight-core processors per node (2GB memory per core) and 3744 Ivy nodes with 2 ten-core processors per node (3.2GB memory per core). The processor speeds of San and Ivy nodes are 2.6 and 2.8 GHz, respectively.

OVERFLOW Parameters

OVERFLOW, which has been developed and advanced by many contributors for over two decades, has several options that can affect the computational performance. In this report default parameters of version 2.2g corresponding to steady state computations are used. The commonly used default parameters that may influence the performance used for all cases in this report are:

- a) Simple time stepping without Newton Sub-Iterations
- b) Central difference option
- c) Spalart-Allmaras turbulence model
- d) Multi-grid option turned-off
- e) Second order accurate algorithm
- f) Variable time step option
- g) Force-Moment computation turned-off
- h) Viscous terms only in the direction normal to the surface

Timing

There are several ways of determining time for parallel jobs. In this report the differences in the time of the day reported in the PBS output file (commonly known as ‘.o’ file or day file) are used.

Saturn V Model

A model of the Saturn V launch vehicle is shown in Fig 1. A baseline grid system totaling 286 million points [8] was used for the performance studies. This grid topology is shown in Figs. 2a to 2c. Part of the near body grids are shown in Fig 2a along with background grid. Figures 2b and 2c show details of the grid around the 5 engines. The near body grids of engines are embedded in background grids. More details regarding this grid topology can be found in Ref. 8.

The topology has 79 NB grids and 4 OB grids. The OB grids were generated externally using OVERGRID. The largest OB grid has 77 million points. The baseline grid size distribution across grid blocks is shown in Fig. 3.

In order to obtain good parallel scalability, OVERFLOW splits grid zones in the baseline grid into smaller sizes and re-packs them in groups for assignment to compute cores. The procedure for grid splitting is given in Appendix A. The number of grid splits depends on the number of cores requested. Figure 4 shows how the number of grid splits increases with the number of cores. Splitting grids adds additional overlapping grid points. Thus, the total number of grid points increases with the number of cores as shown in Fig. 5.




Fig. 1 Saturn V launch vehicle.



Fig. 2a Portions of near body grids (colored) embedded in off-body grid (grey).




Fig. 2b Side view of box grids around engines.




Fig 2c. View of engine grid from bottom.




Fig .3. Baseline grid distribution across blocks.




Fig. 4 Increase in number of grid splits with cores.

Computations were completed using San and Ivy nodes. For San nodes computations were attempted using 128, 256, 512, 1024, 2048, 4096, and 8192 cores. The memory required for the double precision option of OVERFLOW per grid point is about 60 words (480 bytes). The total memory required is about 140GB. The maximum memory available for each core on San nodes is about 2.0 GB. Since only a portion of full memory is available for user, the case with 128 cores was aborted due to lack of memory to accommodate the largest 77million background grid. The case with 8192 cores was also aborted as one of the nodes ran out of memory. This is

attributed to the memory size being inadequate because of the large number of grid splitting performed within one core during preprocessing.

Using Ivy nodes, computations were successfully completed with 200, 400, 800, 1600, 3200, and 6400 cores.

For both San and Ivy nodes, computations were completed for 10 and 60 iterations. The differences in time from the day-files (.o file) for these two runs were used to compute the time required per iteration. This eliminates the overhead time required for starting and ending the jobs.

In order verify the repeatability of timings, 5 cases at different times of the day are run using minimum required 256 cores of San nodes and 200 cores of Ivy nodes. Figure 6 shows the percentage deviation from mean time needed by San and Ivy nodes. Maximum deviation for San nodes is about 1.5% whereas that for Ivy nodes is around 0.6%. However both variations are within reasonable limits to continue computations.




Fig. 5 Increase in grid size with increase in number of cores used.




Fig. 6 Comparison of times between San and Ivy nodes at different run times.

Figure 7 shows times required per step per grid point on San and Ivy nodes. The increase in grid size due to grid splitting associated with increase in use of cores are included in computing time per step per grid point. Both curves show improvement in performance up to about 2000 cores and then flatten.

The clock speed of an Ivy node is 108% that of a San node. Based on interpolated values at 1000 cores, the ratio of time between Ivy and San nodes is 1.18 and does not reflect the fact that Ivy nodes are slightly faster than San nodes.

Figure 8 shows the plots on a log-log scale along with the ideal scaling curve for both San and Ivy nodes. The time axis is normalized with respect to the time required for 256 cores, the minimum used for San nodes. The performances for both nodes are slightly lower than ideal curve and decrease with increase in number of cores used.




Fig.
7.
Co
mpa

rison between San and Ivy times with increasing use of cores.




Fig. 8. Comparison of performances with the ideal scaling curve.

HART II Full Rotorcraft model

A wind tunnel model of the HART II rotorcraft [9] is shown in Fig. 9. A grid with 32 near-body grid blocks and a total of 27M grid points was used. Figure 10 shows the surface grid and Fig. 11 shows the 3-grid system of each blade.




Fig. 9 HART II Rotorcraft Model




Fig. 10 Surface grid for HART II model

The distribution of grid sizes across grid blocks is shown in Fig. 12. The largest grid size of 4.6 million points corresponds to the near-body grid for the blade. For this case, OB grids were generated using OVERFLOW. For 5% chord resolution in the finest OB grid (level 1), OVERFLOW generates 184 OB grids adding 587 million grid points. Table 1 shows the list of

off-body grids generated by OVERFLOW. Grid size growth with increase in the number of cores is shown in Fig. 13.




Fig. 11 Wing grid of HART II




Fig. 12 Number of grid points in near body blocks

Table 1: Grid splitting data from OVERFLOW output.

NEAR-BODY/OFF-BODY GRID LEVEL SUMMARY (before splitting):

	Level	#Grids	First	Last	#Grid-points	(%)
near-body	0	32	1	32	26941183	(4.4)
off-body	1	128	33	160	559851264	(91.2)
off-body	2	14	161	174	20109670	(3.3)
off-body	3	6	175	180	4682860	(0.8)
off-body	4	6	181	186	1414853	(0.2)
off-body	5	6	187	192	530489	(0.1)
off-body	6	6	193	198	235953	(0.0)
off-body	7	6	199	204	130425	(0.0)
off-body	8	6	205	210	97692	(0.0)
off-body	9	6	211	216	78837	(0.0)
total		216	1	216	614073226	(100.0)

Because of the very large grid size, computations using San nodes were only possible using 512, 1024, 4048, and 5120 cores. Below 512 jobs were aborted due to lack of memory to accommodate the largest NB grid with 4.7 million grid points. Beyond 5120 cores, jobs were aborted due to lack of adequate memory required for grid splitting process carried out in a single core. Because of the larger memory available on Ivy nodes, cases could be run using up to 8000 cores and as few as 200 cores. Computations were made for 10 and 60 iterations to determine the time per iteration per grid point.




Fig 13 Increase in grid size with increase in number of cores




Fig. 14 Comparison times between San and Ivy nodes at different run times.

In order to verify the repeatability of timings, 5 cases at different times of the day are run using minimum required 512 cores of San nodes and 200 cores of Ivy nodes. Figure 14 shows the percentage deviation from mean time needed by San and Ivy nodes. Maximum deviation for San nodes is about 2% whereas that for Ivy nodes is around 0.5%. However, both are within reasonable limits to continue computations.

Figure 14 shows a comparison between times required per step per grid point for San and Ivy nodes. The increase in grid size due to grid splitting associated with increase in use of cores is included in computing time per step per grid point. Both curves show improvement in performance up to about 2000 cores and then flatten.

The clock speed of an Ivy node is 108% that of a San node. Based on interpolated values at 1000 cores, the ratio of time between Ivy and San nodes is 0.87, which reflects the fact that Ivy nodes are slightly faster than San nodes.

Figure 16 shows the plots on a log-log scale along with the ideal scaling curve for San and Ivy nodes. The time axis is normalized with respect to time required for 512 cores, the minimum used for San nodes. Both plots show linear scalability close to ideal up to about 1000 cores. Beyond 1000 cores, San nodes show slightly lower and Ivy nodes show slightly higher performance than ideal curve.

Conclusions and Future work

Performance studies are made for version 2.2g of the widely used OVERFLOW CFD code. Timings reported by system outputs are used. Based on runs made at five different times of the day, it is observed that the performances are repeatable within a maximum of 2% variation for San nodes and 0.6% variation for Ivy nodes. Due to the increase in grid sizes with increasing number of cores, for both Saturn V (286 million point grid) and HART II (614 million point

grid) scalability flattens after about 2000 cores. Further studies using two grid topologies of similar grid size may give a better insight as to how performance depends on the type of grid topology.




Fig. 15 Comparison of scaled time between San and Ivy nodes.




Fig. 16 Comparison of performances with the ideal scaling performance.

Acknowledgements

The author thanks Dennis Jespersen of Ames Research Center for providing consultations during this effort. This work is supported by the High End Computing Capability project of NASA.

References

1. Pleiades Super Computer “High End Computing Capability,” “<http://www.nas.nasa.gov/hecc/resources/pleiades.html>”
2. “Contemporary High Performance Computing Systems: From Petascale toward Exascale”, Chapter 12, edited by Jeffrey Vetter, 2013.
3. Nichols, R. and Bunning, P., “User’s Manual for OVERFLOW 2.1,” University of Alabama at Birmingham, Birmingham, AL, 2008.
4. Pandya, S.A, Huang, A., Espitia, A., and Uranga, A., “Computational Assessment of the Boundary Layer Ingesting Nacelle Design of the D8 Aircraft,” 52nd Aerospace Sciences Meeting, AIAA 2014-0907, National Harbor, MD, Jan. 2014.

5. Guruswamy, G. P., "Navier-Stokes based Unsteady Aerodynamic Computations of Launch Vehicles Undergoing Forced Motions," AIAA 2013-1887, 54th AIAA Structural Dynamics Conf., Apr 2013.
6. Saini, S, Chang, J., and Jin, H, "Performance Evaluation of the Intel Sandy Bridge Based NASA Pleiades Using Scientific and Engineering Applications," High Performance Computing Systems. Performance Modeling, Benchmarking and ...simulation, Springer, Switzerland 2014.
7. Chan, W. M., "Developments in Strategies and Software Tools for Overset Structured Grid Generation and Connectivity," AIAA Paper 2011-3051, 20th AIAA Computational Fluid Dynamics Conference, Honolulu, Hawaii, June 2011.
8. Pandya, S., Chan, W., and Kless, J., "Automation of Structured Overset Mesh Generation for Rocket Geometries," AIAA Paper 2009-3993, 19th AIAA Computational Fluid Dynamics Conference, San Antonio, Texas, June 2009.
9. Guruswamy, G. P., "Time-Accurate Aeroelastic Computations of a Full Helicopter Model using the Navier-Stokes Equations," International J. of Aerospace Innovations, Vol. 5, No. 3+4, Dec. 2013, pp. 73-82.

APPENDIX A – Grid Splitting

This section describes the guidelines used in OVERFLOW code for splitting the grid to maximize parallel efficiency.

1. Calculate a target maximum grid size as $1/2$ the total number of points divided by the number of groups (MPI processes). ($1/2$ is ad-hoc, and is chosen to give us some wiggle-room when packing the groups with grids.)
2. Recursively split each grid in two along the longest grid dimension, until the pieces are all less than or equal to the target maximum grid size. There are some limits and restrictions, basically certain boundary conditions and turbulence model regions (see `omisoft/group/split_dirn.F`):
 - edges of less than 29 points are not split
 - axis wraparound directions are not split
 - C-grid or fold-over directions are not split
 - copy-to/copy-from regions are not split
 - uniform inflow conditions are not split
 - Vortex generator vane source models are not split
 - don't split right next to a BC applied to an interior grid surface
 - Baldwin-Lomax viscous-direction regions are not split
 - Baldwin-Barth viscous directions are not split if they are less than $2*29$ points
3. Distribute the grids into the MPI groups in a round-robin fashion, putting the largest un-assigned grid into the least-full group until all grids are assigned to groups.